34.1 A 5.6GHz 64kB Dual-Read Data Cache for the POWER6™ Processor

J. Davis[1], D. Plass[1], P. Bunce[1], Y. Chan[1], A. Pelella[1], R. Joshi[2], A. Chen[1], W. Huott[1], T. Knips[1], P. Patel[1], K. Lo[1], E. Fluhr[3]

[1]IBM, Poughkeepsie, NY
[2]IBM, Yorktown Heights, NY
[3]IBM, Austin, TX

The 16kB SRAM described here is 1 of 4 copies that comprise the 64kB L1 data cache in the POWER6™ microprocessor core. Designed in 65 nm SOI technology [1], it has several features that enable improved cell stability and performance [2]. It is a unidirectional polysilicon dual-supply hierarchical-BL 8-way set-associative design with a 2-stage pipeline supporting 2 independent reads or 1 write per cycle. Each read port, A or B, provides 4B from 1 of 512 locations, whereas the double-bandwidth write operation provides individual control of 8B to 1 of 256 locations. The upper and lower fields of the 8B write data are directed to either of the 4B read-port outputs via a byte steering signal. A WL disable feature enables SRAM "bypass" operations with minimal circuitry and normal timings. Hardware measurements have shown full functionality at 5.6GHz (1.2V, 5°C).

The macro floorplan is shown in Fig. 34.1.1. The memory cells are grouped into 4 vertical sections, each containing 72 active columns of cells and 2 spare columns for use in skip-over replacement; this provides 8 interleaved sets for a 9b byte. The macro also contains upper and lower halves, each comprised of 16 active and 1 spare sub-array (SA). The spare columns and SAs provide 2-dimensional redundancy totaling 6 repair actions.

Read and write addresses are received in the "central core", where an initial level of decoding and multiplexing is provided by a clock-chopped AND-OR. Figure 34.1.2 shows 1 of 2 identical sets of decoding circuitry for each port. The first-level predecoder uses narrow capture pulses (CP-RD and CP-WRT) and a wider Reset_B (RB) pulse to create stretched return-to-zero pulses that are propagated throughout the decoder tree. The narrow capture pulses eliminate the need for an "L1" latch on the address inputs. Address and port gating in read clock-choppers save power by only activating a single SA per enabled read port.

The second-level predecoder generates an SA select (SS) signal and a row-in-SA (RIS) signal. There are 16 RIS signals, 16 SS signals, 1 spare SS signal and a synchronization pulse (SP) that tracks the circuit and wiring of the SS signals. The SP is used to control timing of the column circuits. The second-level SS predecoder incorporates a skip-over redundancy MUX enabling replacement of a defective SA. The word decoder combines an SS and a RIS signal to select a WL, which activates one pass transistor on one side of each cell in a row. During a write, the write address feeds the decode circuitry for both ports, activating both pass transistors for a given row. This scheme permits the use of a traditional word decoder circuit, which can be shared with single-port SRAMs. Clock gating to the first-level predecoder can be used to inhibit RIS signal generation, effectively bypassing the cell writing.

The 0.75μm² (1.5×0.5) butted-junction 6T "thin" memory cell has each of the 2 pass gates connected to separate metal 3 WLs and metal 1 BLs. Single-ended reads are initiated by activating a WL connected to one of the pass gates, and writes are effected when both WLs for a given row are active.

Figure 34.1.3 shows the details of the hierarchical data path. The pass gates of 16 memory cells are dotted together to form a local BL (LBL) pair; a single LBL serving each port. The sensing scheme used in this design presents a very light load to the LBL. A single PMOS provides the pull-up to the LBL_B node. The pull-down is NMOS, controlled by a programmable analog signal (PREN), which can be used for read margin testing. This structure drives an NMOS which is dotted to the metal 2 global BL (GBL).

Each macro-half has its own set of GBLs that are a collection of 9 drain-dotted NMOSs, 1 from each of 8 SA pairs and 1 from the spare SA. During a read, only the upper or lower GBL from a given column of cells will be active for a given port. During the double-bandwidth write, both the upper and lower GBL pairs will be active with the data being written to lower and upper-order bytes, respectively. A transmission gate at the end of each GBL, before the sense PMOS, serves several purposes. By controlling which transmission gate (upper or lower) is active for a given port, and making use of the fact that during a write the same write data is represented on both GBLs, the write data for the upper or lower order bytes can be steered to either output port. Careful timing of the PE_N and PE_P signals prevents very fast accesses from corrupting the "L5" latch before the data from the previous cycle has been transferred to the "L2" output latch. Each transmission gate also forms a leg of a MUX that implements skip-over redundancy.

The final stage in the read path is an integrated 8-input dynamic MUX and latch. The latch is reset high at the beginning of the cycle and then sampled using 1 of 8 chopped clocks (C2S) derived from set select signals received from the lookup and compare functions contained in the set prediction SRAM (SETP). Interleaving the sets within the SA permits the MUX to be local which saves power and wire resources by driving only the 36 MUX-reduced outputs per port to a remote location.

The key elements of the write path are also shown in Fig. 34.1.3. Writing is accomplished by pulling one LBL in the pair low. In order to prevent a false read on the other LBL, low-V_t source-follower NMOSs are used to prevent excessive drooping from its precharged voltage. In addition to the source-followers, the risk of a false read is further mitigated by holding the gate of the corresponding GBL dotted NMOS low. Combining an SS signal with a byte-write (BW) signal creates a local WRTN signal that feeds the source terminals of the LBL write NMOSs in the selected SA. This saves power by only cycling the LBLs in SAs with both an active WL and an active BW. Figures 34.1.4 and 34.1.5 are a high-level block diagram and an associated timing diagram, showing the relationships between the circuits described above.

The design uses dual power supplies, a logic supply in the 0.9 to 1.2V range, and an SRAM power supply approximately 150mV higher. To improve performance, the elevated SRAM supply powers the memory cells, word decoders, second level predecoders, and the write data drivers (which includes the gates of the source-follower NMOSs on the LBLs). Cell stability is enhanced by precharging the BLs to the lower logic voltage, eliminating BL clamping, and ensuring rapid BL decay by limiting LBL load to 16 cells. Figure 34.1.6, summarizes the key attributes of the 16kB data-cache block.

Acknowledgements:
Brad McCredie, Carl Anderson, Brian Curran, Emmanuel Crabbe, Sean Carey, Steve Runyon, Dan Rodko, Tom Chang, Bryan Robbins, Seth Erlebacher, Eric Cheesebrough, Dave Trost, Roger Purvee, the LSU and Technology teams.

References:
[1] E. Leobandung, et al., "High Performance 65nm SOI Technology with Dual Stress Liner and Low Capacitance SRAM Cell," *Symp. on VLSI Technology*, pp. 126-127, June, 2005.
[2] R. Joshi, et al., "Variability Analysis for Sub-100nm PD SOI SRAM Cell," *ESSCIRC Proc.*, pp. 411-415, Sept., 2004.

 1-4244-0079-1/06/$20.00 ©2006 IEEE

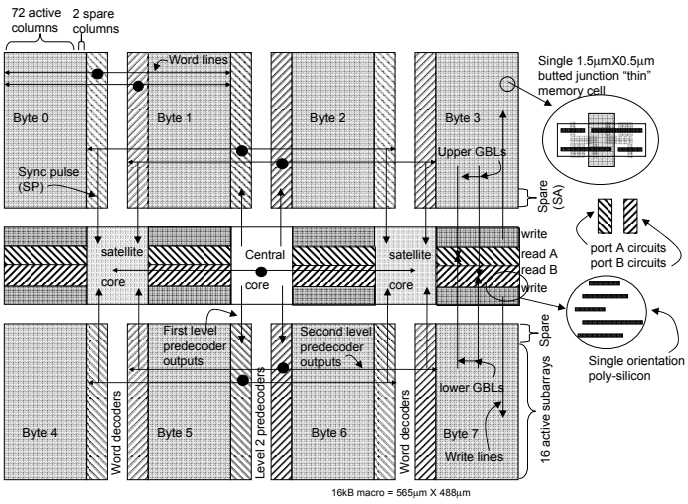

Figure 34.1.1: SRAM floorplan.

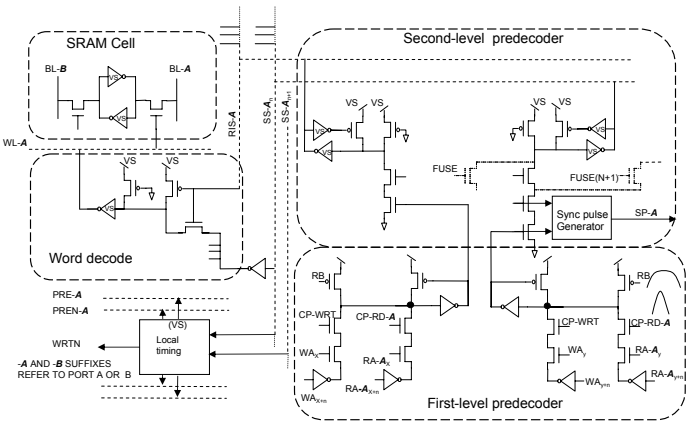

Figure 34.1.2: Address path (Port A).

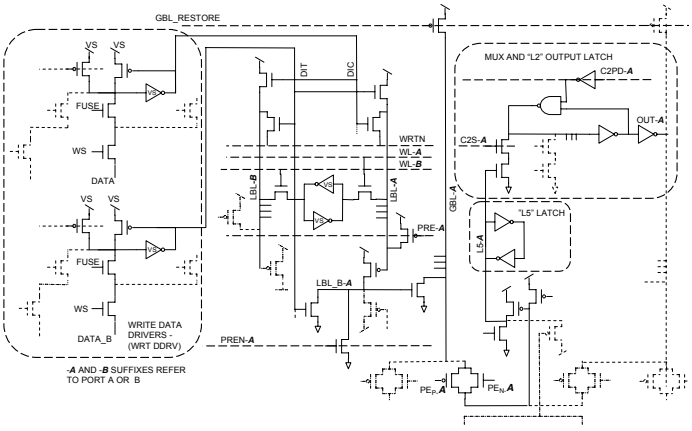

Figure 34.1.3 Data path.

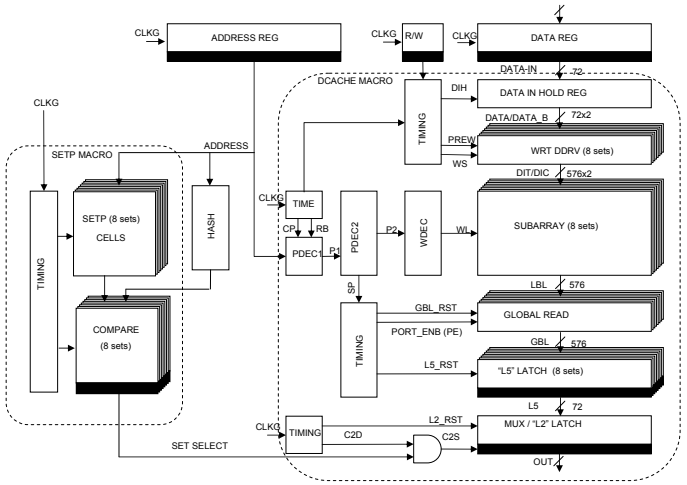

Figure 34.1.4: Circuit relationships.

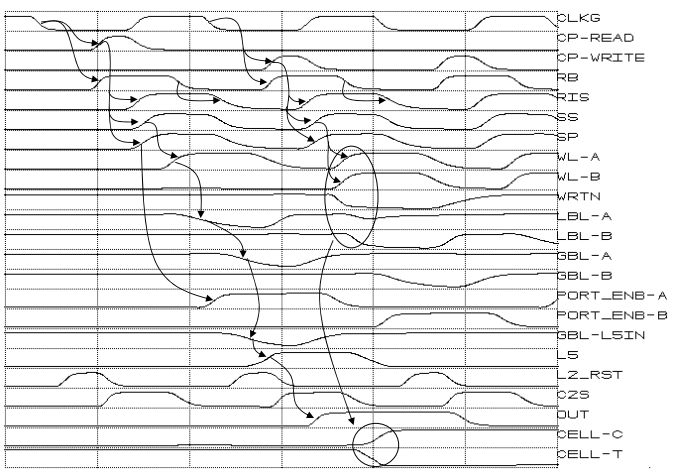

Figure 34.1.5: Simulated waveforms of a read followed by a write.

Macro Frequency	5.6GHz (measured 1.2V V_{LOGIC} and 1.3V V_{SRAM} at 5°C)
Access time	< 200ps (simulated)
Power	250mW (simulated) (50% dual read / 50% writes)
Technology	65nm SOI CMOS
Memory Cell	Split WL 6T butted junction 0.75µm² (1.5 X 0.5) with metal 1 BL and metal 3 WLs
Array Dimensions	Read: 512 entries X 36b (from 1 of 8 sets) X2 ports; Write: 256 entries X 72b (to 1 of 8 sets)
Sensing Scheme	Hierarchical "domino" read with metal 1 local bitlines and metal 2 global bitlines
Area	0.276mm² (565X488)
Efficiency	40% (macro), 58% (subarray)
Redundancy	4 pairs of columns, 2 subarrays
Dual power supply	V_{LOGIC} = 0.9 - 1.2V; V_{SRAM} = V_{LOGIC} +~ .150V
Power saving features	Per port/per subarray clock gating, interleaved sets, high-vt memory cell and non-critical support circuits
Special features	Byte steering, write bypass, byte merge

Figure 34.1.6: Key attributes of a 16kB data cache block.